IN THE SPECIFICATION

Please replace the paragraph beginning at page 1, line 11, with the following replacement paragraph:

---

In a shared-memory multiprocessor system, it appears to a user that all processors read and modify state information in a single shared memory store. A substantial difficulty in implementing such a system, and particularly a distributed version of such a system, is propagating values from one processor to another, in that the actual values are created close to one processor but might be used by many other processors in the system. If the implementation could accurately predict the sharing patterns of a given program, the processor nodes of a distributed multiprocessor system could spend more of their time computing and less of their time waiting for values to be fetched from remote locations. Despite the development of processor features such as non-blocking caches and out-of-order instruction execution, the relatively long access latency in a distributed shared-memory system remains a serious impediment to performance.

A1

---

Please replace the paragraph beginning at page 2, line 22, with the following replacement paragraph:

---

The invention provides improved techniques for determining a set of predicted readers of a data block subject to a write request in a shared-memory multiprocessor system. In accordance with an aspect of the invention, a current set of readers of the data block are determined, and then the set of predicted readers is generated based on the current set of readers and at least one additional set of readers representative of at least a portion of a global history of a directory associated with the data block. In one possible implementation, the set of predicted readers ~~are~~ is generated by applying a function to the current set of readers and one or more additional sets of readers. The function may be, for example, a union function, an intersection function or a pattern-based function, and the directory and data block may be elements of a memory associated with a particular processor node of the multiprocessor system.

A2

---

Please replace the paragraph beginning at page 4, line 20, with the following replacement paragraph:

The invention will be illustrated herein in conjunction with exemplary distributed shared-memory multiprocessor systems. It should be understood, however, that the invention is more generally applicable to any shared-memory multiprocessor system in which it is desirable to provide improved performance through the use of directory-based prediction. The term "multiprocessor system" as used herein is intended to include any device in which retrieved instructions are executed using ~~one~~ two or more processors. Exemplary processors in accordance with the invention may include, for example, microprocessors, central processing units (CPUs), very long instruction word (VLIW) processors, single-issue processors, multi-issue processors, digital signal processors, application-specific integrated circuits (ASICs), personal computers, mainframe computers, network computers, workstations and servers, and other types of data processing devices, as well as portions and combinations of these and other devices.

A3

Please replace the paragraph beginning at page 5, line 3, with the following replacement paragraph:

FIGS. 1 and 2 illustrate the handling of ~~example~~ exemplary read and write requests, respectively, in a distributed shared-memory multiprocessor system 100. The system 100 is an example of one type of system in which the directory-based prediction of the present invention may be implemented. The system 100 includes nodes A, B and C, which are connected to an interconnection network 102 via corresponding network interfaces (NIs) 104A, 104B and 104C, respectively. The nodes A, B and C include processors 106A, 106B and 106C, memories 108A, 108B and 108C, and buses 110A, 110B and 110C, respectively, arranged as shown. Within a given node *i* of the system 100, *i* = A, B, C, the processor 106*i*, memory 108*i* and network interface 104*i* are each coupled to and communicate over the corresponding bus 110*i*.

A4

Please replace the paragraph beginning at page 6, line 27, with the following replacement paragraph:

In the ~~example~~ exemplary implementation of the illustrative embodiment to be described in conjunction with FIG. 4 below, a history depth of four is used, i.e., the predicted set of readers generated for a current write operation on a given block is determined as a function of the current set of readers of that block and the three other most recent sets of readers stored in a predictor shift register.

A5

Please replace the paragraph beginning at page 7, line 3, with the following replacement paragraph:

FIG. 4 shows an example of the operation of a directory-based predictor in the illustrative embodiment of the invention. In this example, a write request is received for a data block X associated with a memory and directory 120. The current readers of the data block X are processors in a set of nodes {a, b, c} of a multiprocessor system which includes nodes denoted a, b, c, d, e, f, g, h, i, j, k, l, m, etc. Each of the nodes may represent a node of a multiprocessor system such as that illustrated in conjunction with FIGS. 1 and 2. The predictor in this example uses a shift register 122 in a manner to be described below.

A6

Please replace the paragraph beginning at page 8, line 7, with the following replacement paragraph:

The choice of union function or intersection function in step 216 of FIG. 5 generally depends on the desired degree of aggressiveness in the data forwarding. For example, in high-bandwidth systems, the more aggressive data forwarding associated with the union function may be more appropriate, while for low-bandwidth systems, the intersection function may be more appropriate. It should be noted that these functions are given by way of example only, and the invention can be implemented using other types of functions. As another example, pattern-based functions can be

A7

used in conjunction with the present invention. Such functions are described in greater detail in, e.g., T. Yeh and Y. Patt, "Two-Level Adaptive Branch Prediction," Proceedings of the 24th Annual ACM/IEEE International Symposium and Workshop on Microarchitecture, Los Alamitos, CA, November 1991, which is incorporated by reference herein.

A1

Please replace the paragraph beginning at page 12, line 21, with the following replacement paragraph:

~~Table~~ Tables 2 and 5 ~~shows~~ show the top ten ~~most sensitive~~ schemes, in terms of specificity, in the set of possible predictors using direct update and forwarded update, respectively. ~~All are union schemes with the maximum history depth used in this example, i.e., a history depth of 4. All schemes are roughly comparable in sensitivity, but with different values of PVP. It is interesting to note that by far the least expensive scheme (union(dir+add$_2$)$^*$) is fifth-best overall in terms of sensitivity.~~

A4

Please replace the paragraph beginning at page 12, line 26, with the following replacement paragraph:

~~Table~~ Tables 3 and 6 ~~shows~~ show the top ten ~~most sensitive~~ schemes, in terms of sensitivity, in the set of possible predictors using direct update and forwarded update, respectively. ~~There is very little difference between the direct- and forwarded-update schemes. Six of the top ten schemes are common to the two lists, and the statistics differ little from column to column.~~ All are union schemes with the maximum history depth used in this example, i.e., a history depth of 4.

A5

Please delete the paragraph beginning at page 13, line 1, as follows:

~~Tables 5 and 6 show the top ten predictors in the set of possible forwarded update predictors in terms of specificity and sensitivity, respectively.~~

Please replace the paragraph of the abstract, at page 19, line 2, with the following rewritten paragraph:

A set of predicted readers are determined for a data block subject to a write request in a shared-memory multiprocessor system by first determining a current set of readers of the data block, and then generating the set of predicted readers based on the current set of readers and at least one additional set of readers representative of at least a portion of a global history of a directory associated with the data block. In one possible implementation, the set of predicted readers are generated by applying a function to the current set of readers and one or more additional sets of readers. ~~The function may be, for example, a union function, an intersection function or a pattern-based function, and the directory and data block may be elements of a memory associated with a particular processor node of the multiprocessor system. The global history of the directory comprises multiple sets of previous readers processed by the directory, with the total number of sets of previous readers corresponding to a designated history depth associated with generation of the set of predicted readers. The prediction process may use additional information in conjunction with the directory information, such as a designated subset of cache address information, processor node identification information, or program counter information.~~